

Risk-Bounded Temporal Logic Control of Continuous-Time Stochastic Systems

Sleiman Safaoui, Lars Lindemann, Iman Shames, Tyler H. Summers

Abstract—Motivated by the recent interest in risk-aware control, we study a continuous-time control synthesis problem to bound the risk that a stochastic linear system violates a given specification. We use risk signal temporal logic as a specification formalism in which distributionally robust risk predicates are considered and equipped with the usual Boolean and temporal operators. Our control approach relies on reformulating these risk predicates as deterministic predicates over mean and covariance states of the system. We then obtain a timed sequence of sets of mean and covariance states from the timed automata representation of the specification. To avoid an explosion in the number of automata states, we propose heuristics to find candidate sequences effectively. To execute and check dynamic feasibility of these sequences, we present a sampled-data control technique based on time discretization and constraint tightening that allows to perform timed transitions while satisfying the continuous-time constraints.

I. INTRODUCTION

The design of safe control laws for autonomous systems has been studied extensively over the past years. For deterministic systems, the safe control synthesis problem is usually cast as a set invariance problem. Proposed solutions consider control barrier functions [1], Hamilton Jacobi reachability analysis [2], or model predictive control [3]. However, when the system is stochastic, e.g., due to uncertainty in the system localization, set invariance has to be interpreted by taking risk into account. Besides safety, the system is subject to performance objectives. As system specification complexity plays a major role in the tractability of the control problem, often simple navigation specifications [4], i.e., going from A to B while avoiding obstacles, or regulation and reference tracking problems are studied. This excludes a large class of specifications such as repetitive specifications (always repeating a certain sequence of events), specifications with strict temporal requirements (reaching some state within a specific time interval then reaching another state), and many others. More complex system specifications have recently been considered using spatio-temporal logics [5], [6]. In this paper, we hence cast the safe control synthesis problem of stochastic system as a risk-aware control synthesis problem with the goal to upper bound the risk that

This work was supported in part by the United States Air Force Office of Scientific Research under Grant FA2386-19-1-4073 and in part by the National Science Foundation under Grant ECCS-2047040.

S. Safaoui and T. Summers are with the School of Engineering at The University of Texas at Dallas, Richardson, TX, USA. E-mail: {sleiman.safaoui, tyler.summers}@utdallas.edu. L. Lindemann is with the School of Engineering at the University of Pennsylvania, Philadelphia, Pennsylvania. Email: larsl@seas.upenn.edu. I. Shames is with the School of Engineering at The Australian National University, Acton, Australia. E-mail: iman.shames@anu.edu.au.

spatio-temporal logic system specifications are violated.

Literature review. Signal temporal logic (STL) is a real-time temporal logic that allows to impose large classes of specifications [7]. Importantly, such specifications permit to define robust semantics that provide information as to what extent a specification is satisfied or violated [8], [9]. The deterministic control synthesis problem has been addressed using optimization techniques [5], [6], [10], machine learning techniques [11]–[13], and automata-based techniques in conjunction with transient control laws [14], [15]. More recently, stochastic control treating the STL specification as a chance constraint have been considered [16]–[18]. These works consider specific notions of risk and assumptions on the state distribution such as Gaussian distributions, and they largely study discrete-time systems. On the other hand, risk-aware control for more simple system specification, i.e., not not complex temporal logic specifications, have been considered in various directions, see e.g., [19]–[22].

In contrast, in this paper we consider risk-aware control under STL specifications. Most closely related to this paper are our previous works [23], [24] in which we consider risk-aware control for STL specifications. While the focus in [23] is on stochastic environments and reactivity, [24] considers stochastic systems and risk, however in a setting where time is discretized and only for the limited fragment of bounded STL specifications. In this paper, we consider a continuous-time stochastic linear system which, to our knowledge, has not been solved with proper formal guarantees. We restrict our attention to linear systems as the distribution's statistics for a nonlinear system are in general hard to estimate.

Contributions. We continue along the lines of our previous work [15] where the continuous-time control synthesis problem for a *deterministic system* under STL specifications is studied. The continuous-time control synthesis problem for a *stochastic system* is more challenging and the state explosion problem can not be addressed as in [15]. First, the problem of finding control laws that achieve timed transitions in the mean and covariance states is difficult. This problem becomes even more difficult when the state distribution is not known, as is the case in this paper. Second, the efficient integration of control laws into the automata representation of the specification is non-trivial. Our contributions are:

- We present, to the best of our knowledge, the first risk-bounded solution to the continuous-time control synthesis problem of stochastic (non-Gaussian) linear systems under STL specifications.
- We propose a sampled-data control technique that performs timed transitions within the space of distributions

and guarantees continuous-time constraint satisfaction.

- As opposed to existing mixed integer linear programming solutions (presented in the deterministic system literature), our method can handle unbounded STL formulas.

II. BACKGROUND

Let \mathbb{R} , $\mathbb{R}_{\geq 0}$, and $\mathbb{Q}_{\geq 0}$ denote the sets of real numbers, non-negative real numbers, and non-negative rational numbers respectively. Let \otimes denote the Kronecker product. Given a matrix, the $vec(\cdot)$ operator stacks its vectors. $diag(\cdot)$ is a diagonal matrix of the arguments. I_n and 0_n are the $n \times n$ identity and zero matrices. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space where Ω is the sample space, \mathcal{F} is a σ -algebra of subsets of Ω , and \mathbb{P} is a probability measure on \mathcal{F} . Given random vectors $X_1, X_2 : \Omega \rightarrow \mathbb{R}^n$, the expected value of X_1 with respect to \mathbb{P} is denoted by $\mathbb{E}[X_1]$ and their covariance is $Cov(X_1, X_2) = \mathbb{E}[X_1 X_2^T] - \mathbb{E}[X_1] \mathbb{E}[X_2]$. We abbreviate positive semi-definite as psd.

A. Real-Time Temporal Logics

Signal interval temporal logic (SITL) is a specification formalism that allows describing a desired system behavior for deterministic systems. A predicate $\mu : \mathbb{R}^n \rightarrow \mathbb{B}$ is a Boolean-valued function that depends on a function $\alpha : \mathbb{R}^n \rightarrow \mathbb{R}$, also referred to as the predicate function. For a given $x \in \mathbb{R}^n$, the predicate is true $\mu(x) = \top$ if $\alpha(x) \geq 0$ and false $\mu(x) = \perp$ if $\alpha(x) < 0$. Let M be a set of atomic predicates $M := \{\mu_1, \dots, \mu_{|M|}\}$. For $\mu \in M$, the SITL syntax is given by

$$\phi ::= \top \mid \mu \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 U_I \phi_2 \quad (1)$$

where ϕ , ϕ_1 , and ϕ_2 are STL formulas and where U_I is the until operator with time interval $I \subseteq \mathbb{Q}_{\geq 0}$ that is not a singleton; \neg and \wedge encode negations and conjunctions. Based on these operators, one can further derive the operators: $\phi_1 \vee \phi_2 := \neg(\neg\phi_1 \wedge \neg\phi_2)$ (disjunction operator), $F_I \phi := \top U_I \phi$ (eventually operator), and $G_I \phi := \neg F_I \neg\phi$ (always operator).

An SITL formula ϕ is evaluated over deterministic signals $x : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$, potential trajectories of a deterministic system. When x satisfies the SITL formula ϕ at time t , we denote this by $(x, t) \models \phi$. The continuous-time STL semantics [8] (define when x satisfies ϕ at t) are inductively defined as:

- $(x, t) \models \mu$ iff $\alpha(x(t)) \geq 0$
- $(x, t) \models \neg\phi$ iff $(x, t) \not\models \phi$
- $(x, t) \models \phi_1 \wedge \phi_2$ iff $(x, t) \models \phi_1$ and $(x, t) \models \phi_2$
- $(x, t) \models \phi_1 U_I \phi_2$ iff $\exists t'' \in t \oplus I$, $(x, t'') \models \phi_2$ and $\forall t' \in (t, t'')$, $(x, t') \models \phi_1$

An SITL formula ϕ is satisfiable if $\exists x$ such that $(x, 0) \models \phi$. Such an SITL formula can be translated into a timed automaton TST_ϕ [15] (see Appendix A for a brief summary). From a timed automaton TST_ϕ , one can obtain plans which can be thought of as requirements on how each predicate $\mu_i(x(t))$ in M and hence the signal $x(t)$ has to evolve over time t (see [15] for details). While SITL is defined over deterministic signals (and hence for deterministic systems), in this paper, we are interested in stochastic systems.

B. Continuous Time Stochastic System

We consider the stochastic linear system described by a stochastic differential equation

$$dX(t) = (AX(t) + Bu(X(t)))dt + dW(t), \quad X(0) = X_0, \quad (2)$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$ are the (constant) dynamics and input matrices. We assume that (A, B) is stabilizable. $X(t) \in \mathbb{R}^n$ is the state, $u : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a state feedback control law, and $X_0 : \Omega \rightarrow \mathbb{R}^n$ is a random variable describing the unknown initial state. $dW(t)$ is differential Brownian motion. The stochastic integral $\int_0^T dW = W(T) - W(0)$ is a possibly non-Gaussian random variable with zero mean $\bar{W} := \mathbb{E}[W(t)] = 0 \ \forall t$, and covariance $\Sigma := \mathbb{E}[W(t)W(t)^T] = diag(\sigma_1^2, \dots, \sigma_n^2)$ where $\sigma_i^2 := Var(W_i(t)) = Cov(W_i(t), W_i(t))$. Often, $W(t)$ is assumed to be Gaussian (some continuous-time stochastic systems literature assumes that without explicitly stating it). However, in this work we *do not* make such assumptions. Instead, to promote robustness to uncertainties in the distribution, we consider a moment based ambiguity set:

$$\begin{aligned} \mathcal{P}^W &:= \{\mathbb{P}^W \mid \mathbb{E}_{\mathbb{P}^W}[W] = \bar{W}, \\ &\quad Cov_{\mathbb{P}^W}(W) = \mathbb{E}_{\mathbb{P}^W}[(W - \bar{W})(W - \bar{W})^T] = \Sigma\} \end{aligned} \quad (3)$$

that is, $W(t)$ can belong to any distribution with mean \bar{W} (in this case $\bar{W} = 0_n$) and psd covariance matrix Σ both of which are assumed to be known.

In this paper, we consider the feedback control law

$$u(X(t)) = K(t)X(t) + k(t), \quad (4)$$

which we motivate in the next section, where $K(t) \in \mathbb{R}^{m \times n}$ is a feedback gain and $k(t) \in \mathbb{R}^m$ is an open-loop control signal. The dynamics in (2) thus become:

$$dX(t) = (A + BK(t))X(t)dt + Bk(t)dt + dW(t). \quad (5)$$

C. Mean and Covariance Dynamics

Consider the dynamics in (5). The state mean is denoted by $\bar{X}(t) := \mathbb{E}[X(t)]$ and the mean dynamics are given by:

$$\dot{\bar{X}}(t) := \mathbb{E}[\dot{X}(t)] = (A + BK(t))\bar{X}(t) + Bk(t), \quad \bar{X}(0) =: x_0 \quad (6)$$

The state covariance is denoted by $P(t) := Cov(X(t), X(t))$ and the covariance dynamics are given by

$$\begin{aligned} \dot{P}(t) &= (A + BK(t))P(t) + P(t)(A + BK(t))^T + \Sigma, \quad P(0) = P_0 \\ \iff \text{vec}(\dot{P}(t)) &= (I_n \otimes (A + BK(t)) + \\ &\quad (A + BK(t)) \otimes I_n) \text{vec}(P(t)) + \text{vec}(\Sigma). \end{aligned} \quad (7)$$

Remark 1. The choice of feedback control law (4) results in having k in the mean dynamics (6) (feedforward term) and K in the covariance dynamics (7) (limit covariance growth).

The mean and covariance dynamics are both given by first order ODEs and can be stacked into a single ODE:

$$\dot{\mathcal{X}}(t) = \mathcal{A}(K(t))\mathcal{X}(t) + \mathcal{B}U(t) \quad (8)$$

where

$$\begin{aligned}\mathcal{A}(K) &:= \begin{bmatrix} A + BK & 0 \\ 0 & (I_n \otimes (A + BK) + (A + BK) \otimes I_n) \end{bmatrix} \\ \mathcal{B} &:= \begin{bmatrix} B & 0 \\ 0 & I_{n \times n} \end{bmatrix}, \quad \mathcal{X}(t) := \begin{bmatrix} \bar{X}(t) \\ \text{vec}(P) \end{bmatrix} \\ \mathcal{U}(t) &:= \begin{bmatrix} k(t) \\ \text{vec}(\Sigma) \end{bmatrix} \quad (\text{vec}(\Sigma) \text{ is a constant}).\end{aligned}$$

We chose to include $\text{vec}(\Sigma)$ (constant) in \mathcal{U} to retain the traditional linear system format. We also assume the following.

Assumption 1. *We assume that there exist closed sets \mathcal{X}, \mathcal{U} and $\mathcal{M} \in \mathbb{R}$ such that $\|\mathcal{X}\| \leq \mathcal{M} \quad \forall (\mathcal{X}, \mathcal{U}) \in \mathcal{X} \times \mathcal{U}$.*

Since (8) is linear, it is Lipschitz continuous. This, and Assumption 1 will become relevant in §IV-B. Regarding the set \mathcal{X} , 1) we consider bounds on the mean state (e.g. physical bounds on the system) so that $\bar{X}_{\min} \leq \bar{X} \leq \bar{X}_{\max}$ where $\bar{X}_{\min}, \bar{X}_{\max} \in \mathbb{R}^n$ are known, and 2) we assume $A + BK(t)$ is stable by the choice of $K(t)$, thus since Σ is psd, P will have a unique symmetric psd steady state value which is the solution to a Lyapunov equation [25, Thm 22]; since the covariance dynamics are linear, $\text{vec}(P)$ will be bounded $0 \leq \text{vec}(P) \leq D \in \mathbb{R}^{n \times n}$ for some D . As for \mathcal{U} , $\text{vec}(\Sigma)$ is a constant and we assume that the open-loop control $k(t)$ is bounded $k_{\min} \leq k \leq k_{\max}$.

III. PROBLEM STATEMENT

In the stochastic framework introduced above, the satisfiability problem of an SITL formula ϕ is ill-posed, i.e., whether or not a stochastic linear system $X(t)$ satisfies the formula ϕ . In fact, an atomic predicate $\mu(X(t))$ becomes a random variable. We hence redefine the atomic predicates as risk predicates. Intuitively, the truth value of a risk predicate $\mu^{Ri}(X)$ is true if the risk of violating the predicate is small.

The risk of violation is obtained by using a risk measure, i.e. a function that maps a random variable to a real number. Let C denote all measurable functions from the sample space Ω to \mathbb{R}^n , i.e., all random variables. Then, a risk measure is defined as $\rho: C \rightarrow \mathbb{R}$. We use the distributionally robust value at risk (DR-VaR) which is a coherent risk metric defined as $\inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}[-\alpha(X) \leq 0] \geq 1 - \eta$ for some risk threshold $\eta \in (0, 1)$. The DR-VaR satisfies certain desirable axioms [26], see [24, §II] for more details on risk measures.

Formally, we define the risk predicate as:

$$\mu^{Ri}(X) := \begin{cases} \top & \text{if } \rho(-\alpha(X)) \leq \eta \\ \perp & \text{otherwise} \end{cases} \quad (9)$$

The risk SITL (RiSITL) syntax is defined as

$$\phi ::= \top \mid \mu^{Ri} \mid \neg \phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 U_I \phi_2. \quad (10)$$

where $\mu^{Ri} \in M^{Ri}$ for the risk predicates $M^{Ri} = \{\mu_1^{Ri}, \dots, \mu_{|M|}^{Ri}\}$, while the other operators have the same meaning as in SITL.

The semantics of RiSITL are different in how the risk predicates are evaluated. Instead of $(x, t) \models \mu$ iff $\alpha(x(t)) \geq 0$ in SITL, we have $(X, t) \models \mu^{Ri}$ iff $\rho(-\alpha(X)) \leq \eta$, while the other operators follow as in the SITL semantics presented

earlier. $(X, t) \models \phi$ indicates that the stochastic linear system with dynamics \dot{X} satisfies the RiSITL formula ϕ at time t . We are now ready to state the formal problem definition.

Problem 1. Given the system (5) and an RiSITL ϕ per (10), find the control variables $K(t)$ and $k(t)$ so that $(X, 0) \models \phi$.

IV. RISK-BOUNDED TEMPORAL LOGIC CONTROL

Our solution to Problem 1 consists of a reformulation of the risk predicates (§IV-A), an optimization-based controller (§IV-B) for timed transitions (see Definition 1 later), and the decomposition of the RiSITL specification into a sequence of timed transitions (§IV-C - §IV-E). In particular, we generate a candidate sequence of timed automaton transitions that we feed into our optimization-based controller to check for dynamic feasibility of this candidate sequence. We note that our solution is sufficient, i.e., sound but not *complete*.

A. Atomic Predicate Reformulation

Consider now linear predicate functions $\alpha(X) = a^T X + b$ where $a \in \mathbb{R}^n$, $b \in \mathbb{R}$. The following reformulation holds

$$\begin{aligned}\rho(-\alpha(X)) \leq \eta &\Leftrightarrow \rho(-(a^T X + b)) \leq \eta \\ &\Leftrightarrow \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}[-a^T X - b \leq 0] \geq 1 - \eta \\ &\Leftrightarrow a^T \bar{X} + b - \underbrace{\sqrt{\frac{1 - \eta}{\eta}} \|P^{1/2} a\|_2}_{:= H} \geq 0 \quad (11)\end{aligned}$$

where the last step follows from [27, Thm 3.1]. The risk predicate (9) is now a *deterministic* risk-tightened predicate:

$$\mu^{Ri}(X) := \begin{cases} \top & \text{if } a^T \bar{X} + b - H \|P^{1/2} a\|_2 \geq 0 \\ \perp & \text{otherwise.} \end{cases} \quad (12)$$

Remark 2. *One way to interpret (12) is to view $\mu^{Ri}(X)$ as a time varying halfspace in the mean dynamics \bar{X} . This follows as $P(t)$ is specified *a priori* when $K(t)$ is chosen in advance. Intuitively, the halfspaces will be tightened so that predicate functions $\alpha(X)$ specifying “goal regions” shrink and “obstacle regions” expand with higher uncertainty.*

B. Control of Timed Transitions

Consider now two polytopes S_1 and S_2 that are subsets of $\mathbb{R}^{n+n \times n}$ and that are either connected or intersecting. We define the timed transition problem as follows.

Definition 1 (Timed Transition). Given polytopes S_1 and S_2 and a transition time $T > 0$, then a feedback gain matrix $K(t)$ and a control law $\mathcal{U}(t)$ achieve a timed transition from $\mathcal{X}(0) \in S_1$ into S_2 at time T if the following is satisfied:

$$\dot{\mathcal{X}}(t) = \mathcal{A}(K(t))\mathcal{X}(t) + \mathcal{B}\mathcal{U}(t) \quad \forall t \in [0, T] \quad (13a)$$

$$\mathcal{U}(t) \in \mathcal{U}, \quad \forall t \in [0, T] \quad (13b)$$

$$\mathcal{X}(t) \in S_1, \quad \forall t \in [0, T] \quad (13c)$$

$$\mathcal{X}(T) \in S_2. \quad (13d)$$

A timed transition under a feedback gain matrix $K(t)$ and a control law $\mathcal{U}(t)$ hence occurs when the mean and the covariance dynamics as well as the inputs bounds \mathcal{U} are

respected, and the system state $\mathcal{X}(t)$ transitions from the set S_1 , in which the system starts, into the set S_2 at time T .

In §IV-C, the sets S_1 and S_2 will encode conjunctions of predicates $\mu^{Ri}(X)$ as per (12). For instance, we may have $S_1 := \{(\bar{X}, \text{vec}(P)) \in \mathbb{R}^{n+n \cdot n} \mid a^T \bar{X} + b - H \left\| P^{1/2} a \right\|_2 \geq 0\}$ for a single predicate function. It is clear that the polytopes are convex in \bar{X} . For example, S_1 may encode a room and S_2 may encode a corridor next to the room. For a robot with initial state in S_1 , a timed transition here would require the robot to stay in the room and transition into the corridor exactly at time T . Towards finding a feedback gain matrix $K(t)$ and a control law $\mathcal{U}(t)$ we propose the following architecture.

1) Finding $K(t)$: To efficiently integrate our control law with the timed automata representation of ϕ (presented in the next section), we propose to first select the closed-loop feedback gain $K(t)$ for a timed transition. Doing so will allow us to determine the covariance matrix $P(t)$ at all times. Without loss of generality, we consider a constant gain K over the entire timed-transition period and require K to satisfy three requirements (related to Assumption 1):

- K must be stabilizing, i.e. $A + BK$ must be stable
- K must keep the dynamics $\dot{\mathcal{X}}$ bounded by \mathcal{M}
- K must keep the covariance bounded by D .

Additionally, K should be chosen to keep \mathcal{M}, D small enough to avoid the problem becoming infeasible when the atomic predicates (12) are tightened. We note that this choice of K may not exist even if a solution to (13) exists. This is one factor that makes our solution methodology only sufficient.

2) Finding $\mathcal{U}(t)$: Once K is determined, the covariance matrix $P(t)$ becomes fully defined for all times $t \geq 0$. We now find a control law $k(t)$, which determines $\mathcal{U}(t)$, to achieve the timed transition as in (13) by solving the optimization problem where open-loop control $k(t)$ is minimized:

$$\min_{k(t)} \int_0^T k(t)^T R k(t) \quad (14a)$$

$$\text{s.t. } (13a) - (13d). \quad (14b)$$

where R is psd. Note that (14) is convex in $k(t)$, but infinite dimensional due to the continuous time formulation resulting in an infinite number of decision variables and constraints. Solving this finite horizon, continuous-time, constrained, open-loop control problem exactly may be possible with Pontryagin's optimality principle, but is difficult in general. However, to facilitate a sampled-data implementation, we solve the above optimization problem by discretizing time as in [28] while tightening the constraints to be able to satisfy the original constraints in continuous time.

Our Assumption 1 is similar to [28, Asm 1]. We note that \mathcal{M} can be found by solving the following problem:

$$\begin{aligned} \mathcal{M} = \arg \max & \quad \|\dot{\mathcal{X}}\| \\ \text{s.t. } & \quad \mathcal{X} \in \mathcal{X}, \quad \mathcal{U} \in \mathcal{U}. \end{aligned}$$

Let us next define the signed distance function $dist_S : \mathbb{R}^{n+n \cdot n} \rightarrow \mathbb{R}$ similar to [28, (5)] as: $dist_S(\mathcal{X}) := \min_{\mathcal{Y} \in S} \|\mathcal{X} - \mathcal{Y}\| - \inf_{\mathcal{Z} \in S^C} \|\mathcal{X} - \mathcal{Z}\|$ i.e., if $\mathcal{X} \in S$, then the distance is the negative distance to the boundary of \mathcal{X} while

if $\mathcal{X} \notin S$ the distance is the positive distance to the boundary. Hence, constraints of the form $\mathcal{X}(t) \in S$ are equivalent to $dist_S(\mathcal{X}(t)) \leq 0$. For simplicity, we opt for a uniform discretization of the time interval $[0, T]$ as $\mathbf{T} := \{\mathcal{T}_0, \dots, \mathcal{T}_N\}$ such that $\mathcal{T}_0 = 0$, $\mathcal{T}_N = T$, and $\delta\mathcal{T} = \frac{T}{N}$ is the time step. Then, using [28] we can guarantee continuous-time satisfaction of the constraints in (13a)-(13d) by tightening them further. With $K(t) = K$ for all $t \in [0, T]$ selected as discussed before and assuming $k(t)$ is a zero-order hold control law, so that we have $k[\mathcal{T}]$ for $\mathcal{T} \in \mathbf{T}$, we propose the following problem to find a controller that satisfies (13) in continuous-time:

$$\min_k \sum_{i \in \{0, \dots, N\}} k[\mathcal{T}_i]^T R k[\mathcal{T}_i] \quad (15a)$$

$$\text{s.t. } \dot{\mathcal{X}}(t) = \mathcal{A}(K)\mathcal{X}(t) + \mathcal{B}\mathcal{U}(t) \quad \forall t \in [0, T] \quad (15b)$$

$$\mathcal{U}(\mathcal{T}) \in \mathcal{U}, \quad \mathcal{T} \in \mathbf{T} \quad (15c)$$

$$dist_{S_1}(\mathcal{X}(\mathcal{T})) \leq -\mathcal{M} \cdot \delta\mathcal{T}, \quad \mathcal{T} \in \mathbf{T} \quad (15d)$$

$$\mathcal{X}(T) \in S_2. \quad (15e)$$

Theorem 1. Given K and \mathcal{M} as in §IV-B.1. Let $k[\mathcal{T}]$ be a solution to (15). Then the constraints of the timed transition problem (14) are satisfied, i.e. a timed transition from S_1 into S_2 as per Definition 1 can be achieved under K and $k[\mathcal{T}]$.

Proof. The proof follows by applying [28, Theorem 1] with the augmented dynamics (8).

Remark 3. The dynamics (15b) are deterministic linear dynamics in \mathcal{X} with a closed-form solution at the sampled times given by: $\mathcal{X}(\mathcal{T}) = e^{\mathcal{A}(K)\mathcal{T}}\mathcal{X}(0) + \int_0^{\mathcal{T}} (e^{\mathcal{A}(K)(\mathcal{T}-t)}Bk(t))dt$ where the interval can be split into several integrals because $k(t) = k[\mathcal{T}_i]$ for $t \in [\mathcal{T}_i, \mathcal{T}_{i+1}]$ [25].

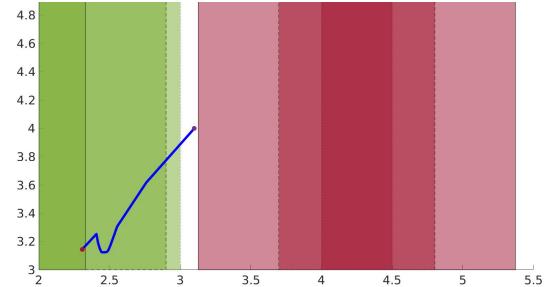


Fig. 1. Constraints: not tightened (dotted), DR-tightening (dashed), and DR- and DT-tightening (solid) for a goal region (left, green) and obstacle region (right, red). A timed transition (blue) moves the robot into the goal.

Example 1. To demonstrate the different tightening procedures described, consider three atomic predicate functions in 2D $\alpha_1 := [-1, 0]X + 3 \geq 0$, $\alpha_2 := [-1, 0]X + 4 \geq 0$, and $\alpha_3 := [1, 0]X - 4.5 \geq 0$. The first defines a goal region (reach) and the other two define an obstacle region (avoid). We use a risk bound $\eta = 0.5$ for $\mu_1^{Ri}(X)$ and $\eta = 0.1$ for $\mu_2^{Ri}(X)$ and $\mu_3^{Ri}(X)$ which are the risk-tightened predicates as in (12). For single integrator dynamics with $K(t) = \text{diag}(-5, -5) \forall t$ and $\delta\mathcal{T} = 0.01$, we have $\mathcal{M} \cdot \delta\mathcal{T} = 0.5712$. With $\Sigma = \text{diag}(0.1, 0.1)$, $P(0) = \mathbf{0}_2$, and $T = 2$, we get $P(2) = \text{diag}(0.0101, 0.0101)$. In Fig. 1, we plot the (not tightened) constraints, the risk-tightened (DR-tightened) constraints at

$t = 2$ and the risk and discrete-time-tightened (DR- and DT-tightened) constraints at $t = 2$ with dotted, dashed, and solid edges respectively. Notice how the goal ‘‘shrinks’’ while the obstacle ‘‘expands’’ as $P(t)$ increases. The timed transition (blue) is the solution to (15) that moves the robot into the goal (from the purple to the red point). The discrete-time points stay between the DR- and DT-tightened predicates of the goal and obstacle for all $t \in \{\mathcal{T}_0, \dots, \mathcal{T}_{N-1}\}$, thus satisfying (15d). At $t = T$, the agent satisfies the DR-tightened predicates (15e).

C. Risk-Based Automaton

As mentioned before, avoiding the state explosion problem for stochastic systems is challenging and can not be addressed as in [15]. This is particularly the case as the efficient integration of control laws for stochastic systems into the timed-automata representation of the specification is non-trivial. In our proposed problem solution, we first translate the RiSITL specification ϕ into a timed automaton TST_ϕ similarly to [15], but with risk predicates μ^{Ri} instead of predicates μ . The procedure closely follows Appendix A, but we will next briefly describe the difference for RiSITL and refer to the appendix for more intuition. We abstract $\phi(M^{Ri})$ into a MITL specification $\varphi(P)$ (see Appendix B for an introduction to MITL) where P is a set of propositions that replaces the set of risk predicates M^{Ri} . The MITL formula φ is then translated into a timed automaton TST_ϕ according to [29]. Each state s and transition δ in TST_ϕ now encode intersections of constraints of the form $\{(\bar{X}, \text{vec}(P)) \in \mathbb{R}^{n+n \cdot n} \mid a^T \bar{X} + b - H \left\| P^{1/2} a \right\|_2 \geq 0\}$ which we denote by $\lambda(s)$ and $\lambda(\delta)$. We then perform the following operations on TST_ϕ to obtain TST_ϕ :

- [$\bar{O1}$] Remove any s in TST_ϕ if $\nexists (\bar{X}, P), (\bar{X}, P) \models \lambda(s)$.
- [$\bar{O2}$] Remove each transition δ in TST_ϕ if $\nexists (\bar{X}, P)$ such that $(\bar{X}, P) \models \lambda(\delta)$.

Operations $\bar{O1}$ and $\bar{O2}$ are taken without the consideration of the system dynamics (8). With respect to Remark 2, we view $P(t)$ as apriori given when $K(t)$ is chosen in advance so that the constraints $\lambda(s)$ and $\lambda(\delta)$ can be seen as time-varying in the mean state. We can hence, instead, perform operations $\bar{O1}$ and $\bar{O2}$ for a fixed P and obtain a tightened automaton denoted by TST_ϕ^{tight} . Some choices of P are:

- $P = 0_n$: [$\bar{O1}$], [$\bar{O2}$] become [$O1$], [$O2$] from [15]
- $P \in \{P_{\max,i} \mid \forall i \in \{1, \dots, |M|\}\}$ such that $P_{\max,i} = \arg \sup_P \left\| P^{1/2} a_i \right\|_2$ where $0 \leq \text{vec}(P) \leq D$: Here, every predicate (12) is tightened by the maximum amount of its $\left\| P^{1/2} a \right\|_2$ term (we call this *maximum tightening*).

In these cases, [$\bar{O1}$] and [$\bar{O2}$] can be constructed as simple feasibility problems as described in [30].

Example 2. Consider the illustrative example in Fig. 2. Ignoring the environment bounds, the goal and obstacle regions are defined with eight atomic predicates $\mu_1^{Ri}, \dots, \mu_8^{Ri}$. The atomic predicates are indicated with dashed lines and the halfspace where the predicate is satisfied is indicated by an arrow (top figure). There are two dark green goal

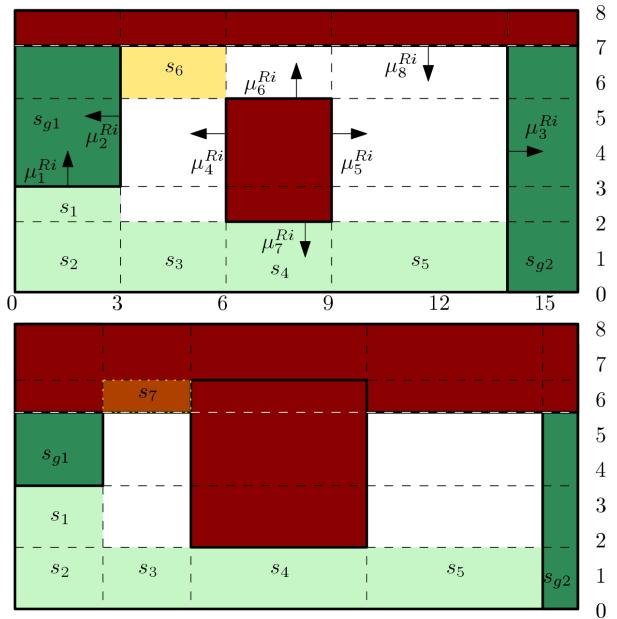


Fig. 2. Automaton states changing as P changes

regions: 1) left $g_1 := \mu_1^{Ri} \wedge \mu_2^{Ri}$ and 2) right $g_2 := \mu_3^{Ri}$. Their atomic predicates point inwards. There are two red obstacle regions: 1) middle $o_1 := \neg \mu_4^{Ri} \wedge \neg \mu_5^{Ri} \wedge \neg \mu_6^{Ri} \wedge \neg \mu_7^{Ri}$ and 2) top $o_2 := \neg \mu_8^{Ri}$. Their atomic predicates point outwards. The top figure is for $P(0) = 0_2$. The bottom is for some steady state $P(\infty) > P(0)$ (its value is irrelevant here). Some examples of automaton states are: $s_1 := \neg \mu_1^{Ri} \wedge \mu_2^{Ri} \wedge \neg \mu_3^{Ri} \wedge \mu_4^{Ri} \wedge \neg \mu_5^{Ri} \wedge \neg \mu_6^{Ri} \wedge \neg \mu_7^{Ri} \wedge \mu_8^{Ri}$ and $s_7 := \mu_1^{Ri} \wedge \neg \mu_2^{Ri} \wedge \neg \mu_3^{Ri} \wedge \mu_4^{Ri} \wedge \neg \mu_5^{Ri} \wedge \mu_6^{Ri} \wedge \neg \mu_7^{Ri} \wedge \neg \mu_8^{Ri}$. Notice that with $P(0)$ (top), s_6 is non-empty, but with $P(\infty)$ (bottom) it is empty and removed by [$\bar{O1}$]. The reverse happens for s_7 . The states s_1, s_2, s_3, s_4, s_5 are non-empty and remain connected for all tightening values $P(0) \leq P(t) \leq P(\infty)$.

D. Risk-Based Control: General Case

Towards an efficient solution, we propose to find candidate sequences of timed transitions from TST_ϕ^{tight} and post-hoc check for dynamic feasibility by means of the control technique derived in §IV-B. The tightened automaton TST_ϕ^{tight} helps us to guide the search process to find a feasible solution and to decrease the search space.

One approach is to consider the maximum tightening case described earlier. This is the most conservative approach and results in the most robust solution. [$O1$] and [$O2$] are applied to TST_ϕ with every atomic predicate maximally tightened (e.g. Fig. 2 (bottom)). The resulting automaton is TST_ϕ^{tight} .

From the tightened automaton TST_ϕ^{tight} , we find a sequence of timed automaton transitions using graph search techniques (see [15], [31] for details). These automaton transitions are defined by alternating transitions between automaton states (in zero time) and transitions within the same automaton state (in finite time). This returns a sequence of the form, $(s_0, 0) \xrightarrow{\delta_0} (s_1, 0) \xrightarrow{\tau_1} (s_1, \tau_1) \xrightarrow{\delta_2} (s_2, \tau_1) \xrightarrow{\tau_2} (s_2, \tau_1 + \tau_2) \xrightarrow{\delta_3} (s_3, \tau_1 + \tau_2)$. $(s_0, 0) \xrightarrow{\delta_0} (s_1, 0)$ is an example of a transition

between automaton states, and $(s_1, 0) \xrightarrow{\tau_1} (s_1, \tau_1)$ is a transition within the same state. The former transition occurs instantaneously; it is the exact moment the switch happens from one automaton state to another (e.g. the exact moment a robot reaches the boundary between two rooms). The latter transition occurs in the same automaton state while the dynamics evolve in time (e.g. the robot moves in one room to reach the boundary to the next room). To simplify notation we use: $(s_1, 0) \xrightarrow{\tau_1, \delta_2} (s_2, \tau_1) := (s_1, 0) \xrightarrow{\tau_1} (s_1, \tau_1) \xrightarrow{\delta_2} (s_2, \tau_1)$.

The timed transitions of Definition 1 accomplish $(s_i, \tau_i) \xrightarrow{\tau, \delta} (s_j, \tau_i + \tau)$. For example, consider $(s_1, 0) \xrightarrow{\tau_1, \delta_2} (s_2, \tau_1)$. The automaton state s_1 corresponds to a convex set of mean states that satisfy its input label $\lambda(s_1)$, i.e. $S_1 := \lambda(s_1)$. This and τ_1 being the transition time T are handled by (15d). Similarly, s_2 corresponds to the mean states of $S_2 := \lambda(s_2)$ (hence (15e)).

Algorithm 1: Risk-Based Timed Transition Control

```

input : Candidate automaton transition sequence
output:  $\mathcal{U}(t)$  or Failure
1 Choose  $K(t) = K$  for this transition;
2  $s = s_0$ ,  $t = 0$ ;
3 for  $\tau, \delta$  in automaton transition sequence such that
    $(s, t) \xrightarrow{\tau, \delta} (s', t + \tau)$  do
4   solved,  $k(\mathcal{T}) = \text{solve\_timed\_transition}(s, \tau, \delta)$ ;
5   if not solved then
6     break;
7   Store  $k(t)$ ;
8    $s = s_{\text{next}}$ ,  $t = t + \tau$ ;

```

Then, for a timed automaton transitions sequence from $TST_{\phi}^{\text{tight}}$, we use Algorithm 1 to check if it is dynamically feasible. For every $(s_i, \tau_i) \xrightarrow{\tau, \delta} (s_j, \tau_i + \tau)$, we use §IV-B.1 and §IV-B.2 to find K and $k(t)$ (lines 1, 4). If we fail to find them, we stop and try a different sequence of timed automaton transitions. If all transitions succeed, then we have found $u(t)$ and Problem 1 is solved. Note that this heuristic approach is *only sufficient* to finding a solution.

Theorem 2. Consider the automaton $TST_{\phi}^{\text{tight}}$ obtained by performing operations $[\overline{O1}]$ and $[\overline{O2}]$ for some fixed matrix P , e.g., the maximal tightening. If a sequence of timed automaton transitions $(s_i, \tau_i) \xrightarrow{\tau, \delta} (s_j, \tau_i + \tau)$ can be found such that the corresponding timed transitions (Definition 1) can be achieved by K and $k(t)$ as per (15), then the linear stochastic system $X(t)$ satisfies the RiSITL formula ϕ , i.e. $(X, 0) \models \phi$ and Problem 1 is solved.

Proof. Ignoring the dynamics, the existence of a sequence of timed automaton transitions ensures that the formula ϕ is satisfiable [15, Lemma 2]. From Theorem 1, if K and $k(t)$ are found for an automaton transition, then the resulting trajectory is dynamically feasible and satisfies the continuous-time risk-based constraints. Applying this to every automaton transition completes the proof. \square

Remark 4. In practice, when solving (15) repeatedly in Algorithm 1, the DT-tightening may sometimes render (15) infeasible. A simple workaround, which works well for small $\delta\mathcal{T}$, is to remove this DT-tightening term $\mathcal{M} \cdot \delta\mathcal{T}$ from (15d) for a few initial and final sampled times (e.g. $\mathcal{T}_0, \mathcal{T}_1, \mathcal{T}_2$ and $\mathcal{T}_{N-2}, \mathcal{T}_{N-1}, \mathcal{T}_N$).

E. Risk-Based Control: Special Case

We consider a special case which allows for a less-conservative solution. While the method can be generalized to \mathbb{R}^n , we will discuss the \mathbb{R}^2 case only. Consider $a \in \{[0 \pm 1]^T, [\pm 1 0]^T\}$. Thus, an automaton state s is a 2D rectangular area. We thus have $a^T P(t) a \in \{\pm P(t)_{(1,1)}, \pm P(t)_{(2,2)}\}$ (i.e. \pm the first or second diagonal elements of $P(t)$) and hence $\|P^{1/2} a\|_2 = \sqrt{a^T P a} \in \{\sqrt{P_{(1,1)}}, \sqrt{P_{(2,2)}}\}$. This makes it easier to describe the conditions on $P(t)$ for which an automaton state s is non-empty (i.e. for which there exists (\bar{X}, P) such that $(\bar{X}, P) \models \lambda(s)$). Example 3 illustrates these conditions.

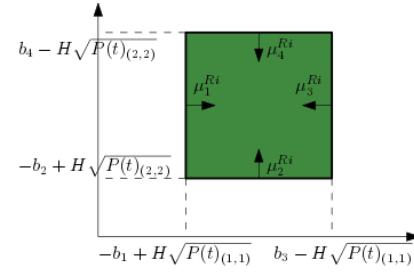


Fig. 3. Special case P conditions example

Example 3. Consider $\mu_1^{Ri} \wedge \mu_2^{Ri} \wedge \mu_3^{Ri} \wedge \mu_4^{Ri}$ in Fig. 3. For the rectangular set to be non-empty, the distances between parallel edges must be positive (note $P_{(1,1)}, P_{(2,2)} \geq 0$, P is psd):

- $-b_1 + H\sqrt{P_{(1,1)}} \leq b_3 - H\sqrt{P_{(1,1)}} \iff \left(\frac{b_1 + b_3}{2}\right)^2 \geq P_{(1,1)}$
- $-b_2 + H\sqrt{P_{(2,2)}} \leq b_4 - H\sqrt{P_{(2,2)}} \iff \left(\frac{b_2 + b_4}{2}\right)^2 \geq P_{(2,2)}$.

We now introduce operation $[\overline{O3}]$ to augment the states of the automaton TST_{ϕ} with conditions on $P(t)$.

$[\overline{O3}]$ For all s in TST_{ϕ} , add the two conditions on $P(t)_{(1,1)}$ and $P(t)_{(2,2)}$ (a la Example 3) for which s is non-empty.

After $[\overline{O3}]$, the resulting TST_{ϕ} has the same states but with extra guards on the states to guide the planning search. $[\overline{O3}]$ can be applied to TST_{ϕ} after $[\overline{O1}], [\overline{O2}]$. Then, we can follow the same procedure described in §IV-D.

V. NUMERICAL EXAMPLE

Consider the environment of Figure 2 (top) and a robot with dynamics $A = \text{diag}(0.07, 0.1)$, $B = \text{diag}(1, 1)$. We use $k_{\max} = -k_{\min} = 30$, $\delta\mathcal{T} = 0.01$, $\Sigma = \text{diag}(0.1, 0.1)$, $K = \text{diag}(-0.9, -0.5)$ for all transitions, and risk bound $\eta = 0.1$ for all constraints except the environment bounds where we use $\eta_{\text{env}} = 0.4$. For the specification, consider the repetitive task $\phi = G_{[0,3+23\omega]}(\neg o_1 \wedge \neg o_2) \wedge F_{[0,3]}g_1 \wedge G_{[3,3+23\omega]}(F_{[0,11]}g_2 \wedge F_{[11,23]}g_1)$ where ω is the number of repetitions of going from g_1 to g_2 and back (g_1, g_2, o_1, o_2 were defined in Example 2).

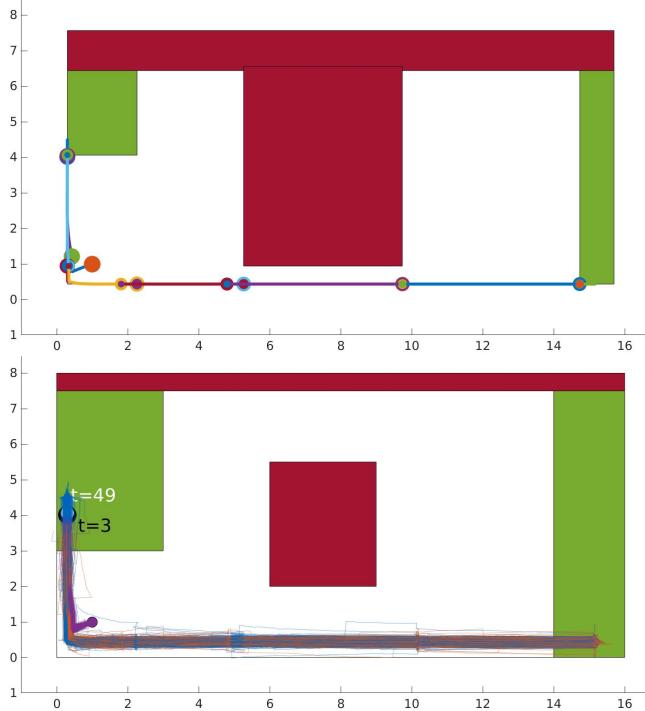


Fig. 4. (Top) Timed transitions of a solution for the considered example. Circles mark the start and end of transition; their radius decreases as time increases. (Bottom) 10000 Monte Carlo rollouts for the given solution. Trajectories: (purple) from the initial condition to left goal, (blue) from the left goal the to right goal, (orange) from the right goal to the left goal.

Using the approach in §IV-E, with the robot starting at $x_0 = [1, 1]$ with $P(0) = 0_2$ we get the sequence of timed automaton transitions: $(s_2, 0) \xrightarrow{\tau_1=1, \delta_1} (s_1, 1) \xrightarrow{\tau_2=2, \delta_2} (s_{g1}, 3) \xrightarrow{\tau_3=1, \delta_3} (s_1, 4) \xrightarrow{\tau_4=2, \delta_4} (s_2, 6) \xrightarrow{\tau_5=2, \delta_5} (s_3, 8) \xrightarrow{\tau_6=2, \delta_6} (s_4, 10) \xrightarrow{\tau_7=2, \delta_7} (s_5, 12) \xrightarrow{\tau_8=2, \delta_8} (s_{g2}, 14) \xrightarrow{\tau_9=2, \delta_9} (s_5, 16) \xrightarrow{\tau_{10}=2, \delta_{10}} (s_4, 18) \xrightarrow{\tau_{11}=2, \delta_{11}} (s_3, 20) \xrightarrow{\tau_{12}=2, \delta_{12}} (s_2, 22) \xrightarrow{\tau_{13}=2, \delta_{13}} (s_1, 24) \xrightarrow{\tau_{14}=2, \delta_{14}} (s_{g1}, 26)]^\omega$ where $[\cdot]^\omega$ represents repeating the transitions ω -times (the robot starts in s_2 ; all these automaton states are marked in Fig. 2). We use (15) to find $k(t)$ per Algorithm 1. The resulting sequence of timed transitions, for $\omega = 2$, is shown in Fig. 4 (top) with the constraints tightened using the steady-state covariance $P(\infty) = \text{diag}(0.0602, 0.125)$. The repeated transitions appear superimposed since they are almost identical. The trajectory starts at $(1, 1)$, moves up to the first goal, then repetitively visits the second goal then the first goal. We run 10000 Monte Carlo simulations. The noise is sampled from a 3 degree of freedom, 0-mean, Σ -covariance, student-t distribution. The rollouts are plotted in the original environment in Fig. 4 (bottom). The black and white circle in g_1 are reached at $t = 3$ and $t = 49$ sec respectively. The robot never collided with the obstacles, but violated the environment bounds in 0.3% $\ll \eta_{env} = 40\%$ of the cases.

VI. CONCLUSION

We present a risk-bounded controller for continuous-time stochastic, non-Gaussian, linear system under STL specifi-

cations. In particular, we use RiSITL to specify constraints rooted in axiomatic risk theory and reformulate DR-VaR constraints into deterministic risk-tightened constraints. Then, we consider timed transitions and tighten our constraints further to account for the discrete-time implementation of a sampled data system without loosing continuous-time guarantees. From there, we show how these timed transitions can be used to verify the dynamically-feasibility of timed automaton transitions from a risk-based automaton.

APPENDIX

A. SITL to Timed Signal Transducers

An SITL formula ϕ can be translated into a language equivalent timed signal transducer, i.e., a timed automaton [15]. We will need some of the machinery presented in [15] despite working with RiSITL, and hence summarize the translation from SITL to timed signal transducer. The first step is to abstract the SITL formula ϕ into an MITL formula φ (see Appendix B for a description of MITL). We use $\phi(M)$ to make explicit that the SITL formula ϕ depends on the set of predicates M . We abstract the SITL formula $\phi(M)$ into an MITL formula $\varphi(P)$ essentially by replacing predicates M in $\phi(M)$ by a set of propositions P . For $i \in \{1, \dots, |M|\}$, associate with each $\mu_i \in M$ a proposition p_i and let $P := \{p_1, \dots, p_{|M|}\}$. Let then $\varphi(P) = \phi(P)$.

The translation from MITL to timed signal transducer mainly follows [29]. Let $c \in \mathbb{R}_{\geq 0}^O$ be a vector denoting O clocks. These clocks can be reset by the reset function $r : \mathbb{R}_{\geq 0}^O \rightarrow \mathbb{R}_{\geq 0}^O$. Clocks evolve with time when visiting a state of a timed signal transducer, while clocks may be reset during transitions between states. We define clock constraints as Boolean combinations of conditions of the form $c_o \leq k$ and $c_o \geq k$ for some $k \in \mathbb{Q}_{\geq 0}$. Let $\Phi(c)$ denote the set of all clock constraints over clock variables in c .

Definition 2 (Timed Signal Transducer [29]). A timed signal transducer is a tuple $TST := (\mathcal{S}, s_0, \Lambda, \Gamma, c, \iota, \Delta, \lambda, \gamma, \mathcal{A})$ where \mathcal{S} is a finite set of states, s_0 is the initial state with $s_0 \cap \mathcal{S} = \emptyset$, Λ and Γ are a finite sets of input and output variables, respectively, $\iota : \mathcal{S} \rightarrow \Phi(c)$ assigns clock constraints over c to each state, Δ is a transition relation so that $\delta = (s, g, r, s') \in \Delta$ indicates a transition from $s \in \mathcal{S} \cup s_0$ to $s' \in \mathcal{S}$ satisfying the guard constraint $g \subseteq \Phi(c)$ and resetting the clocks according to r ; $\lambda : \mathcal{S} \cup \Delta \rightarrow BC(\Lambda)$ and $\gamma : \mathcal{S} \cup \Delta \rightarrow BC(\Gamma)$ are input and output labeling functions where $BC(\Lambda)$ and $BC(\Gamma)$ denote the sets of all Boolean combinations over Λ and Γ , respectively, and $\mathcal{A} \subseteq 2^{\mathcal{S} \cup \Delta}$ is a generalized Büchi acceptance condition.

To construct a timed signal transducer TST_φ that encodes the MITL formula φ , we follow the algorithm presented in [29]. For a signal $d : \mathbb{R}_{\geq 0} \rightarrow \mathbb{B}^{|P|}$, it holds that $(d, 0) \models \varphi$ if and only if d satisfies the generalized Büchi acceptance condition of TST_φ (see [29] for a definition). As the MITL formula φ is an abstraction of the SITL formula ϕ , we perform two operations on TST_φ to obtain the timed signal transducer TST_ϕ encoding ϕ , as presented in [15]:

[O1] Remove any $s \in \mathcal{S}$ if $\nexists x \in \mathbb{R}^n$ so that $x \models \lambda(s)$. Remove the corresponding s from \mathcal{A} .

[O2] Remove any $\delta := (s, g, r, s') \in \Delta$ if $\nexists x \in \mathbb{R}^n$ so that $x \models \lambda(\delta)$. Remove the corresponding δ from \mathcal{A} .

The modified TST_ϕ is denoted by $TST_\phi := (\mathcal{S}^\phi, s_0, \Lambda, \Gamma, c, \iota, \Delta^\phi, \lambda, \gamma, \mathcal{A}^\phi)$. For a signal $x : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$, it holds that $(x, 0) \models \phi$ if and only if x satisfies the generalized Büchi acceptance condition of TST_ϕ .

B. Metric Interval Temporal Logic (MITL)

While system specifications are given as SITL formulas, an intermediate step is needed when one wants to obtain a timed signal transducer that encodes the SITL formula. The difference between SITL and MITL is that SITL considers predicates, while MITL considers propositions. Instead of $\mu(x)$, MITL only considers propositions p where $p = \top$ if the proposition holds and $p = \perp$ if the proposition does not hold. Propositions are hence abstractions of predicates so that MITL can be seen as an abstraction of SITL. Let P be a set of propositions. For $p \in P$, the MITL syntax is: $\varphi ::= \top \mid \perp \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 U_I \varphi_2$ with a similar interpretation of the operators as for SITL. Note that we use φ to denote MITL formulas, while we use ϕ to denote SITL formulas.

MITL semantics are similar to SITL semantics. An MITL formula φ is interpreted over a Boolean signal $d : \mathbb{R}_{\geq 0} \rightarrow \mathbb{B}^{|P|}$ that corresponds to truth values of the propositions in P over time. Define the projection of d onto $p \in P$ as $\text{proj}_p(d) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{B}$. The only difference of the MITL semantics compared to the SITL semantics is now that instead of $(x, t) \models \mu$ iff $\alpha(x(t)) \geq 0$, we have $(d, t) \models p$ iff $\text{proj}_p(d)(t) = \top$, while the other operators follows as in the SITL semantics [29, Sec. 4]. The expression $(d, t) \models \varphi$ indicates that d satisfies the MITL formula φ at time t .

REFERENCES

- [1] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [2] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin, “Fastrack: A modular framework for fast and guaranteed safe motion planning,” in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 1517–1522.
- [3] H. Chen and F. Allgöwer, “A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability,” *Automatica*, vol. 34, no. 10, pp. 1205–1217, 1998.
- [4] D. V. Dimarogonas, S. G. Loizou, K. J. Kyriakopoulos, and M. M. Zavlanos, “A feedback stabilization and collision avoidance scheme for multiple independent non-point agents,” *Automatica*, vol. 42, no. 2, pp. 229–243, 2006.
- [5] V. Raman, A. Donzé, M. Maasoumy, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia, “Model predictive control with signal temporal logic specifications,” in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 81–87.
- [6] S. Sadreddini and C. Belta, “Robust temporal logic model predictive control,” in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2015, pp. 772–779.
- [7] O. Maler and D. Nickovic, “Monitoring temporal properties of continuous signals,” in *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*. Springer, 2004, pp. 152–166.
- [8] A. Donzé and O. Maler, “Robust satisfaction of temporal logic over real-valued signals,” in *Proc. Int. Conf. FORMATS*, Klosterneuburg, Austria, September 2010, pp. 92–106.
- [9] G. E. Fainekos and G. J. Pappas, “Robustness of temporal logic specifications for continuous-time signals,” *Theoretical Computer Science*, vol. 410, no. 42, pp. 4262–4291, 2009.
- [10] N. Mehdiipour, C.-I. Vasile, and C. Belta, “Arithmetic-geometric mean robustness for control from signal temporal logic specifications,” in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 1690–1695.
- [11] A. G. Puranic, J. V. Deshmukh, and S. Nikolaidis, “Learning from demonstrations using signal temporal logic in stochastic and continuous domains,” *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6250–6257, 2021.
- [12] P. Varnai and D. V. Dimarogonas, “On robustness metrics for learning stl tasks,” in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 5394–5399.
- [13] W. Liu and C. Belta, “Model-based safe policy search from signal temporal logic specifications using recurrent neural networks,” *arXiv preprint arXiv:2103.15938*, 2021.
- [14] D. Gundana and H. Kress-Gazit, “Event-based signal temporal logic synthesis for single and multi-robot tasks,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3687–3694, 2021.
- [15] L. Lindemann and D. V. Dimarogonas, “Efficient automata-based planning and control under spatio-temporal logic specifications,” in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 4707–4714.
- [16] S. S. Farahani, R. Majumdar, V. S. Prabhu, and S. Soudjani, “Shrinking horizon model predictive control with signal temporal logic constraints under stochastic disturbances,” *IEEE Transactions on Automatic Control*, vol. 64, no. 8, pp. 3324–3331, 2018.
- [17] D. Sadigh and A. Kapoor, “Safe control under uncertainty with probabilistic signal temporal logic,” 2016.
- [18] P. Jagtap, S. Soudjani, and M. Zamani, “Formal synthesis of stochastic systems via control barrier certificates,” *IEEE Transactions on Automatic Control*, 2020.
- [19] S. Singh, Y. Chow, A. Majumdar, and M. Pavone, “A framework for time-consistent, risk-sensitive model predictive control: Theory and algorithms,” *IEEE Trans. Autom. Control*, vol. 64, no. 7, pp. 2905–2912, 2018.
- [20] M. P. Chapman, J. Lacotte, A. Tamar, D. Lee, K. M. Smith, V. Cheng, J. F. Fisac, S. Jha, M. Pavone, and C. J. Tomlin, “A risk-sensitive finite-time reachability approach for safety of stochastic dynamic systems,” in *Proc. Am. Control Conf.*, Philadelphia, PA, June 2019, pp. 2958–2963.
- [21] M. Ahmadi, M. Ono, M. D. Ingham, R. M. Murray, and A. D. Ames, “Risk-averse planning under uncertainty,” in *Proc. Am. Control Conf.*, Denver, CO, July 2020, pp. 3305–3312.
- [22] M. Schuurmans and P. Patrinos, “Learning-based distributionally robust model predictive control of markovian switching systems with guaranteed stability and recursive feasibility,” in *Proc. Conf. Decis. Control*, Jeju Island, Republic of Korea, December 2020, pp. 4287–4292.
- [23] L. Lindemann, G. J. Pappas, and D. V. Dimarogonas, “Reactive and risk-aware control for signal temporal logic,” *IEEE Transactions on Automatic Control*, 2021.
- [24] S. Safaoui, L. Lindemann, D. V. Dimarogonas, I. Shames, and T. H. Summers, “Control design for risk-based signal temporal logic specifications,” *IEEE Control Systems Letters*, vol. 4, no. 4, pp. 1000–1005, 2020.
- [25] D. Simon, *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons, 2006.
- [26] S. Zymler, D. Kuhn, and B. Rustem, “Worst-case value at risk of nonlinear portfolios,” *Management Science*, vol. 59, no. 1, pp. 172–188, 2013.
- [27] G. C. Calafiore and L. El Ghaoui, “On distributionally robust chance-constrained linear programs,” *Journal of Optimization Theory and Applications*, vol. 130, no. 1, pp. 1–22, 2006.
- [28] F. A. Fontes and L. T. Paiva, “Guaranteed constraint satisfaction in continuous-time control problems,” *IEEE control systems letters*, vol. 3, no. 1, pp. 13–18, 2018.
- [29] T. Ferrère, O. Maler, D. Ničković, and A. Pnueli, “From real-time logic to timed automata,” *Journal of the ACM (JACM)*, vol. 66, no. 3, p. 19, 2019.
- [30] A. Bemporad and M. Morari, “Control of systems integrating logic, dynamics, and constraints,” *Automatica*, vol. 35, no. 3, pp. 407–427, 1999.
- [31] R. Alur and T. A. Henzinger, “The benefits of relaxing punctuality,” *Journal of the ACM*, vol. 43, no. 1, pp. 116–146, 1996.